

Exploiting Calibration-Induced Artifacts in Lossless Compression of Hyperspectral Imagery

Aaron B. Kiely, *Senior Member, IEEE*, and Matthew A. Klimesh, *Member, IEEE*

Abstract—Algorithms for compression of hyperspectral data are commonly evaluated on a readily available collection of Airborne Visible/Infrared Imaging Spectrometer (AVIRIS) images. These images are the end product of processing raw data from the instrument, and their sample value distributions contain artificial regularities that are introduced by the conversion of raw data values to radiance units. It is shown that some of the best reported lossless compression results for these images are achieved by algorithms that significantly exploit these artifacts. This fact has not been widely reported and may not be widely recognized. Compression performance comparisons involving such algorithms and these standard AVIRIS images can be misleading if they are extrapolated to images that lack such artifacts, such as unprocessed hyperspectral images. In fact, two of these algorithms are shown to achieve rather unremarkable compression performance on a set of more recent AVIRIS images that do not contain appreciable calibration-induced artifacts. This newer set of AVIRIS images, which contains both calibrated and raw images, is made available for compression experiments. To underscore the potential impact of exploiting calibration-induced artifacts in the standard AVIRIS datasets, a compression algorithm is presented that achieves noticeably smaller compressed sizes for these datasets than is reported for any other algorithm.

Index Terms—AVIRIS, hyperspectral imagery, lossless data compression, predictive compression.

I. INTRODUCTION

The enormous data volumes produced by hyperspectral imagers have sparked an interest in specially designed data compression techniques to increase data return from spaceborne hyperspectral imagers as well as to provide more efficient ground archival and distribution of hyperspectral imagery. Scientist end-users of such imagery are often reluctant to accept any distortion introduced by lossy compression, and thus much of the focus has been on lossless compression methods, e.g., [1]–[11].

Algorithms for compression of hyperspectral data are commonly evaluated on a readily available collection of Airborne Visible/Infrared Imaging Spectrometer (AVIRIS) images. These images have undergone a radiometric calibration procedure that introduces artificial regularities into the data. In this paper we review two of the lossless compression

algorithms that are among the best performing on this data and show that they benefit from this calibration-induced structure, thus potentially creating a misleading impression of how they will compare with other algorithms when applied to, e.g., raw hyperspectral imagery on a spacecraft. To underscore the potential impact of exploiting calibration-induced structure, a compression algorithm is presented that is deliberately designed to exploit this structure and that achieves significantly smaller compressed sizes for the standard AVIRIS images than is reported for any other algorithm. Finally, we present lossless compression results on these images and on some additional hyperspectral images which we make publicly available.

II. THE STANDARD 1997 AVIRIS IMAGES

The calibrated radiance images from the 1997 AVIRIS sample data products¹ are the most widely used data for benchmarking hyperspectral image compression algorithms, both lossless (e.g., [1]–[11]) and lossy (e.g., recently [12]–[16]). These images are listed in Table I. Each image contains 614 samples/line and includes 224 spectral bands covering wavelengths from 370 nm to 2500 nm. Radiance values are stored as 16-bit signed integers, but the raw AVIRIS data from which they were produced consists of 12-bit unsigned values [17].

TABLE I
THE STANDARD 1997 AVIRIS IMAGES.

Site	Features	Lines
Moffett Field	vegetation, urban, water	2031
Jasper Ridge	vegetation	2586
Cuprite	geological features	2206
Lunar Lake	calibration	1431
Low Altitude	high spatial resolution	3689

The radiance images are the product of radiometric calibration, which converts raw digital number (DN) values acquired by the instrument to convenient multiples of radiance units [17], [18]. Of most significance for our interests is the fact that during the calibration procedure the DN values are multiplied by a band- and image-dependent factor that is somewhat larger than one — it is always greater than 1.8 for these images. The effect of this scaling can be seen in the histograms of sample values in individual calibrated bands; it is manifested as regularly spaced local peaks, with the spacing of the peaks essentially equal to the scaling factor. It is primarily this

Submitted October 10, 2008; revised December 12, 2008 and January 26, 2009. This work was supported by the National Aeronautics and Space Administration.

The authors are with the Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109 USA (e-mail: Aaron.B.Kiely@jpl.nasa.gov; Matthew.A.Klimesh@jpl.nasa.gov).

This work has been submitted to the IEEE for possible publication. Copyright may be transferred without notice, after which this version will be superseded.

¹Available at <http://aviris.jpl.nasa.gov/html/aviris.freedata.html>. This set is generally referred to as the 1997 images, although the Low Altitude image was in fact acquired in 1996.

regularity that can be exploited by a compression algorithm. We describe some further details of the calibration procedure below, but understanding the finer details of this procedure is not essential to appreciating the central point of this paper.

The band scaling factors are the product of a radiometric calibration coefficient determined from laboratory calibration measurements, an onboard calibration coefficient derived from an onboard calibrator signal, and an arbitrary factor to scale the radiance units appropriately for representation with 16-bit integers [17]. Figure 1 shows the resulting band scaling factors for two representative images from the standard set.

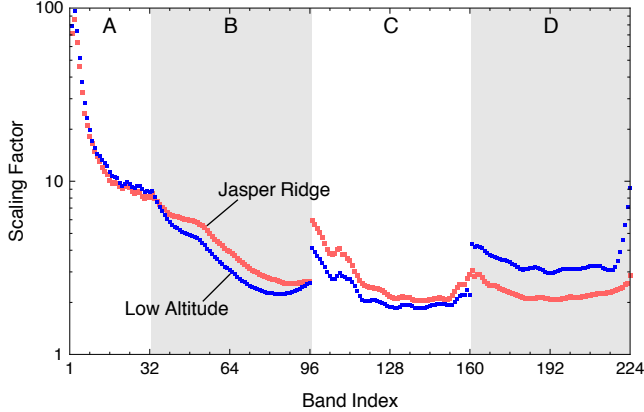


Fig. 1. Band scaling factors between radiance and raw DN values for the Jasper Ridge and Low Altitude images. Scaling factors for the other 1997 images are nearly indistinguishable from those for the Jasper Ridge image. The shading indicates the range of bands corresponding to the four spectrometers (labeled A, B, C, D) used in the AVIRIS instrument [17].

Figure 2 shows a portion of the histogram for band 54 of the Jasper Ridge image. This band has a scaling factor of 5.4, and the local peaks at this spacing are apparent.

Calibration involves more than simply uniformly scaling DN values for each band; additional steps include subtraction of an estimated dark signal level (which varies from line to line) [17], [18] and spectrometer-dependent nonlinear corrections [19]. These additional steps have a tendency to smear the histograms, making the peaks less prominent. But if the standard deviation of the combined adjustment from these corrections is small compared to the scaling factor, then local peaks may still be noticeable; this is apparently the situation in Figure 2.

The histogram regularities are more pronounced for radiance values taken from the same line, presumably because the same dark signal estimate is subtracted from each value in a line. A single line from one band does not contain enough samples to provide a good demonstration of this, but we can produce a histogram for all lines that have a given dark signal value. We have done this for band 54 of the Jasper Ridge image using the most common dark signal value. Figure 3 shows the resulting histogram; we see that the peaks are much more pronounced than those in Figure 2.

III. EXPLOITING CALIBRATION-INDUCED ARTIFACTS

Most of the lossless hyperspectral image compression algorithms in the literature, and all of the ones we discuss in

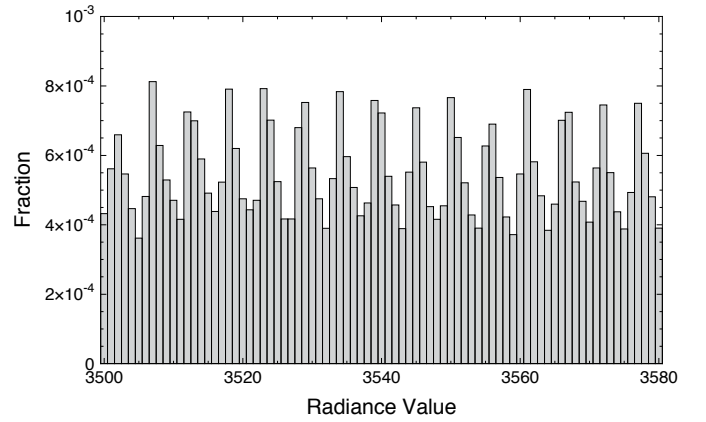


Fig. 2. Portion of histogram of calibrated radiance values in band 54 of the Jasper Ridge image.

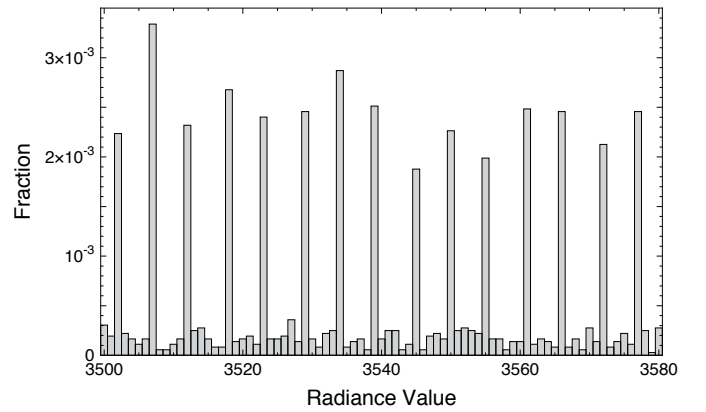


Fig. 3. Portion of histogram of calibrated radiance values in band 54 of the Jasper Ridge image, from the lines with the most common dark signal value.

this article, perform predictive compression. For our purposes, a predictive compression algorithm is defined as follows. If $s_{b,y,x}$ denotes the value of the sample in band b at line y and column x of a hyperspectral image, then to encode the value of $s_{b,y,x}$, a predictive compression algorithm uses the values of previously encoded samples to compute a predicted sample value $\hat{s}_{b,y,x}$. The prediction residual (the difference between $s_{b,y,x}$ and $\hat{s}_{b,y,x}$) is then losslessly encoded using an entropy coder. For convenience, we will also use the notation $s_b(i)$ to denote the i th sample in band b when samples within the band are arranged in raster-scan order.

Under the predictive compression framework, the compressor's entropy coder implicitly or explicitly relies on a probability distribution model to encode prediction residuals. For some predictors, the distribution of prediction residuals may inherit structure such as regularly spaced local peaks from the underlying data. In this case, the compressor may be taking advantage of artifacts in the data when the entropy coder is capable of exploiting such structure, as is typically the case when adaptive arithmetic coding or range coding is used. In the remainder of this section we describe the predictors used by certain compressors and demonstrate that two of them facilitate exploitation of calibration-induced data structure in

this manner.

To our knowledge, the best lossless compression results on the 1997 AVIRIS images are obtained by the locally averaged interband scaling quantized lookup table (LAIS-QLUT) compressor [1], which is a modification of the locally averaged interband scaling lookup table (LAIS-LUT) compressor [2], which in turn is an extension of the lookup table (LUT) compressor presented in [3].

The LUT algorithm [3] predicts the value of sample $s_b(i)$ by finding the most recent index j of the sample in band $b-1$ that matches the value of $s_{b-1}(i)$. That is, the predictor finds the largest $j < i$ such that $s_{b-1}(j) = s_{b-1}(i)$. If such a match exists, then the predictor is $\hat{s}_b(i) = s_b(j)$, otherwise the predictor is set to the value of the sample in the previous band, $\hat{s}_b(i) = s_{b-1}(i)$. A lookup-table facilitates the implementation of this prediction method. Prediction residuals are encoded using adaptive range coding.

The LAIS-LUT refinement of this approach [2] determines the indices j_1, j_2 of the two most recent matches for $s_{b-1}(i)$ in band $b-1$, i.e., the algorithm finds $j_1 < j_2 < i$ such that $s_{b-1}(j_1) = s_{b-1}(j_2) = s_{b-1}(i)$. If there is only one such index j_1 , then the predictor is $\hat{s}_b(i) = s_b(j_1)$. If no match is found, a local interband predictor $\hat{s}'_{b,y,x}$ is computed by scaling the value of $s_{b-1,y,x}$ by the average of band ratios at three neighboring spatial locations:

$$\hat{s}'_{b,y,x} = \frac{1}{3} \left(\frac{s_{b,y-1,x}}{s_{b-1,y-1,x}} + \frac{s_{b,y,x-1}}{s_{b-1,y,x-1}} + \frac{s_{b,y-1,x-1}}{s_{b-1,y-1,x-1}} \right) s_{b-1,y,x}$$

and the predictor is $\hat{s}_b(i) = \text{round}(\hat{s}'_{b,y,x})$. If two matches are found, then $\hat{s}_b(i)$ is equal to one of the two candidates $s_b(j_1), s_b(j_2)$, whichever is closer to $\hat{s}'_{b,y,x}$. Prediction residuals are encoded using adaptive arithmetic coding.

The LAIS-QLUT algorithm [1] modifies LAIS-LUT by uniformly quantizing the value of $s_{b-1}(i)$ used in the lookup table, and by using a different definition of the local interband predictor $\hat{s}'_{b,y,x}$. There are two variations of the LAIS-QLUT algorithm that differ in the method used to select the quantization step size used for each band.

The LUT-based compressors make limited use of spatial correlation, and they exploit spectral dependency using only a single band. And yet, for the 1997 AVIRIS images, on average the LAIS-LUT and LAIS-QLUT compressors outperform substantially more complex compressors such as clustered differential pulse code modulation (C-DPCM) [4]. The C-DPCM compressor partitions spectral vectors into 16 clusters and then applies a separate least-squares optimized 20th order linear predictor to each cluster of each band.

In our analysis and experiments we focus on the LUT and LAIS-LUT methods, as the LAIS-QLUT method appeared in the literature only after we had performed much of this work. In any case, the LAIS-QLUT compressor contains the key features that the LUT and LAIS-LUT compressors use to benefit from regularities in calibrated imagery, so it is reasonable to assume that extending our experiments to the LAIS-QLUT method would produce results similar to those for the LUT and LAIS-LUT methods.

Histograms of prediction residuals resulting from applying the LUT and LAIS-LUT methods to the 1997 AVIRIS images clearly show structure inherited from the calibration-induced structure in the data (Figure 4). The adaptive entropy coding techniques used in the LUT and LAIS-LUT methods are able to exploit this structure for improved compression performance. This phenomenon does not seem to be well-recognized, although it is noted in [5] for the LUT method. For comparison, the figure also shows prediction residuals for the “fast lossless” (FL) compressor of [6], [7]. The FL predictor combines spatial and spectral prediction via adaptive linear filtering, making predictions using previously coded samples in a small neighborhood from the current and three preceding bands.

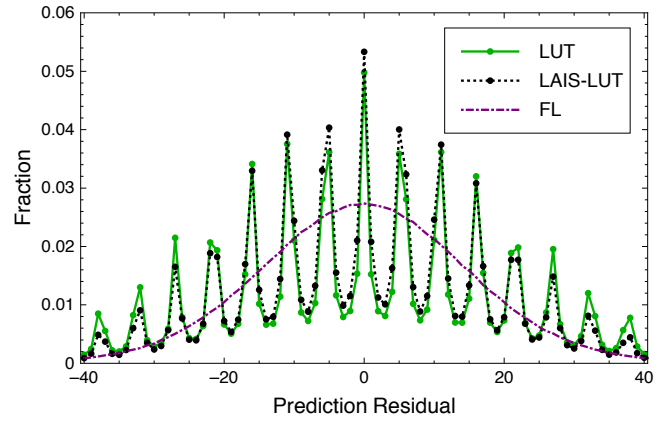


Fig. 4. Histogram of prediction residuals for band 54 of the Jasper Ridge image for LUT, LAIS-LUT, and FL predictors.

The structure in the LUT and LAIS-LUT prediction residual distributions takes the form of somewhat regularly spaced sharp spikes. This structure arises from similarly spaced spikes in the sample distributions (introduced by radiometric calibration) combined with the fact that the LUT and LAIS-LUT predictors usually pick previously encoded values, often from the current line of the image, as predictions.

By contrast, the FL predictor tends to produce prediction residuals with a relatively smooth distribution. Furthermore, the FL compressor uses a simpler entropy coding approach (specifically, adaptive Golomb coding similar to that described in [20]) that would not be particularly effective at encoding residuals with a spiky distribution in any case.

IV. A NEW COMPRESSOR

We now describe a new predictive compressor that deliberately exploits structure introduced by radiometric calibration. The effectiveness of this compressor on the standard 1997 AVIRIS images is significantly better than that of the LUT-based compressors (and all other compressors that we know of), but the primary reason for creating it is to demonstrate the extent to which the calibration-induced structure can be exploited. We refer to the approach as two-stage predictive (TSP) compression. This approach is mildly reminiscent of a simple technique for exploiting sparse histograms in standard

(two-dimensional) images that is described in Section V-B of [20].

A. Prediction

Our TSP compressor makes use of a conventional predictor to form an initial integer-valued prediction $\hat{p}_b(i)$. This initial predictor is intended to exploit spatial and/or spectral correlations in the data but not calibration-induced structure. In our experiments we use the FL predictor for this purpose, but other choices should also work well. To form the ultimate prediction, $\hat{s}_b(i)$, statistics from the initial prediction residuals are weighted by an adaptively updated function that takes advantage of calibration-induced data structure.

During compression, for each band we tabulate counts of the past values of the difference between the initial prediction and the actual sample value; specifically, we let $f_b(r)$ be equal to the number of times this difference has been equal to r in band b . Note that f_b is thus also implicitly a function of the current sample index i . We would typically expect an initial prediction residual close to zero to be most common; f_b should tend to smoothly decay away from its peak, which should be near zero.

We also maintain an adaptively updated weight function $w_b(m)$ defined over all possible sample values m in the band. The weight function is intended to capture information about data regularities introduced by the calibration process. We consider two choices for weight functions. Our first weight function (W1) simply maintains a tally of the past sample values in the band, so that at sample index i we have $w_b(m) = |\{j < i : s_b(j) = m\}|$. The second weight function (W2) gives much higher weight to more recent sample values by assigning $w_b(m) = 2^{-a}$, where a is the number of lines since the value m previously appeared in the band (e.g., $a = 0$ if the sample value m appeared previously in the current line being encoded) or a is the index of the line being encoded if m has not yet appeared. Our W1 definition is motivated by a desire to exploit the sort of overall histogram structure seen in Figure 2, while our W2 definition, which gives highest weight to samples within the same line, is motivated by the more significant local structure suggested by Figure 3. Other weight functions may give better compression performance.

The final predicted value is

$$\hat{s}_b(i) = \arg \max_m f_b(m - \hat{p}_b(i)) \cdot w_b(m).$$

The difference between $\hat{s}_b(i)$ and $s_b(i)$ is losslessly encoded using adaptive arithmetic coding, as we describe below.

As an example illustrating our prediction approach, Figure 5 shows the counts $f_b(m - \hat{p}_b(i))$, the weight function $w_b(m)$, and the product $f_b(m - \hat{p}_b(i)) \cdot w_b(m)$ when encoding a sample near the middle of band 98 in the Jasper Ridge image using weight function W1.

B. Entropy Coding

For our entropy coding stage, we use an arithmetic coding approach that is effective at encoding a source with an arbitrary distribution. Specifically, for each spectral band we apply a

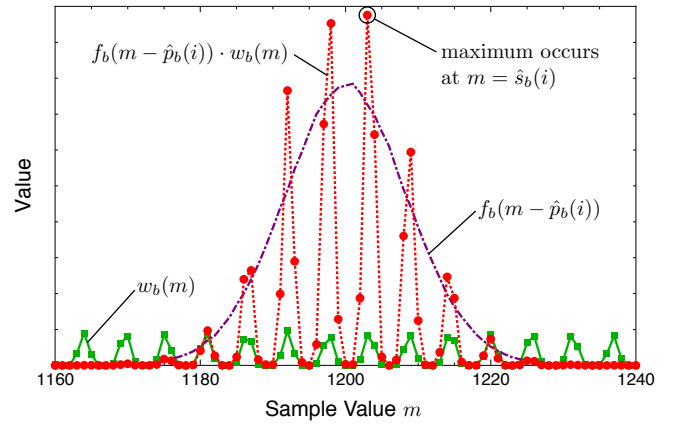


Fig. 5. Example of the functions used to compute the TSP-W1 predictor in band 98 of the Jasper Ridge image. The vertical scales are irrelevant in the computation. The product function is scaled differently for clarity of presentation.

separate instance of the adaptive arithmetic coder of [21] using the publicly available software implementation.² We use the software's adaptive zero-order integer-based model encoding method, setting all parameters to their default values.

We also computed compression results when this same entropy coding approach is applied to prediction residuals produced by the LUT, LAIS-LUT, and FL predictors. We do this in part for convenience, but also because using the same entropy coder can facilitate a fair comparison between different predictors. In our results we designate these modified compressors by appending the # symbol following the algorithm name (e.g., LUT#).

The adaptive zero-order integer-based model method of the arithmetic coding software of [21] cannot operate on prediction residuals directly; instead it takes as input a sequence known as the *pattern* [22], [23] of the residual sequence. In the pattern, the integer n corresponds to the n th new symbol to appear in the residual sequence. Thus each integer in the pattern must be in the range $1, \dots, r + 1$, where r is the largest integer seen in the pattern so far. In order to recover the original residual sequence, the decoder needs to know the correspondence between the residuals and the pattern integers. We encoded this additional information using a somewhat ad hoc technique; since this information accounted for less than 0.7% of the total bit rate on each of the images considered in this paper, improvements to this technique would have little impact on the results and we omit further details.

V. RESULTS

A. 1997 AVIRIS Images

Table II lists the compressed data rates in bits/sample for several lossless compression methods applied to the standard 1997 AVIRIS images. As a reference point, we include results from [8] for the JPEG-LS image compression standard [20] applied independently to spectral bands. The table also includes results from the literature for C-DPCM [4], spectral fuzzy-matching pursuits (S-FMP) [5], LUT [3], LAIS-LUT [2], and

²http://www.cs.mu.oz.au/~alistair/arith_coder/

TABLE II
COMPRESSION PERFORMANCE (RATE IN BITS/SAMPLE) ON THE 1997 AVIRIS IMAGES.

Image	JPEG-LS	SLSQ-OPT	FL	FL#	C-DPCM	S-FMP	LUT	LUT#	LAIS-LUT	LAIS-LUT#	LAIS-QLUT-OPT	TSP-W1	TSP-W2
Moffett Field	8.04	4.99	4.99	4.93	4.62	4.63	5.05	5.05	4.76	4.79	4.62	4.67	4.12
Jasper Ridge	8.38	4.96	4.95	4.87	4.62	4.63	4.95	4.93	4.68	4.70	4.61	4.51	4.08
Cuprite	7.66	4.94	4.91	4.82	4.68	4.66	4.65	4.66	4.47	4.48	4.29	4.53	3.77
Lunar Lake	7.48	4.95	4.91	4.83	4.75	4.66	4.71	4.71	4.53	4.53	4.34	4.36	3.81
Low Altitude	8.00	5.27	5.26	5.18	—	—	—	5.20	—	5.00	—	5.00	4.31
Average (first four images)	7.89	4.96	4.94	4.86	4.67	4.64	4.84	4.84	4.61	4.63	4.47	4.52	3.95
Average (all images)	7.91	5.02	5.00	4.93	—	—	—	4.91	—	4.70	—	4.62	4.02

LAIS-QLUT-OPT [1], which is the LAIS-QLUT variation that selects the optimum lookup table index quantization step size for each band. Results for spectral-oriented least squares using the optimum selection between interband and intraband prediction (SLSQ-OPT) are obtained using data provided by the authors of [9]. We obtained FL results after partitioning images into 512-line segments (with a smaller final segment), and so they are slightly better than results reported in [6], which used 32-line segments. We obtained results for LUT# and LAIS-LUT# using our own implementation of the predictors. The last two columns of the table list compression results for our TSP method using the two different choices of weight function. Note that results for C-DPCM [4] are obtained by compressing images in 512-line segments, omitting some of the lines at the end of each image. All other results in Table II are for complete images.

It can be concluded from comparing the TSP results with the FL# results that both versions of the TSP compressor are gaining an advantage by exploiting calibration-induced artifacts in the standard AVIRIS images. Furthermore, the TSP-W2 compressor provides roughly half a bit per sample improvement over the previous best results.

Figure 6 shows the compressed data rate for each band of the Jasper Ridge image for both versions of the TSP compressor, as well as for the LUT#, LAIS-LUT#, and FL# compressors. Compared to the FL predictor, the TSP predictors and the LUT-based predictors provide the most significant gains in the first 32 bands, where the band scaling factor is largest (cf. Figure 1).

B. AVIRIS Images from the CCSDS Test Set

Calibrated and uncalibrated AVIRIS images have been provided to the Consultative Committee for Space Data Systems (CCSDS) Multispectral and Hyperspectral Data Compression working group for compression testing and evaluation. These images are now available for download.³ This newer AVIRIS data, summarized in Table III, consists of five calibrated and corresponding 16-bit raw images acquired over Yellowstone, WY in 2006, as well as two additional 12-bit uncalibrated images. Each image is a 512-line scene containing 224 spectral bands. We partitioned contiguous portions of the complete runs into 512-line scenes and selected a manageable subset for inclusion in the CCSDS test set. Our aim in selecting

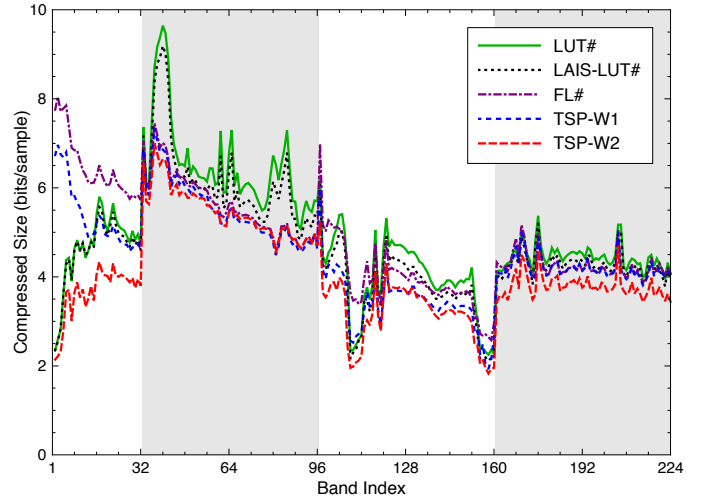


Fig. 6. Compressed data rate for each band of the Jasper Ridge image for the LUT#, LAIS-LUT#, FL#, TSP-W1, and TSP-W2 compressors.

the portions of runs and scenes within the runs was to eliminate non-imaging lines at the start and end of a run, and avoid scenes containing lines affected by an intermittent data corruption anomaly.

TABLE III
AVIRIS IMAGES INCLUDED IN THE CCSDS TEST SET.

Site	Scene Numbers	Year	Samples/Line	Bit Depth	Type
Yellowstone	0, 3, 10, 11, 18	2006	677	16	cal.
Yellowstone	0, 3, 10, 11, 18	2006	680	16	uncal.
Hawaii	1	2001	614	12	uncal.
Maine	10	2003	680	12	uncal.

Table IV gives compression performance results for the FL, FL#, LUT#, LAIS-LUT#, and TSP compressors on the AVIRIS images in the CCSDS test set. The authors of [5] speculate that an algorithm such as LUT that exploits sparse histograms may lose its performance advantage when applied to uncalibrated data. Table IV supports this; even without arithmetic coding, the FL compressor noticeably outperforms the two LUT-based compressors examined, and the TSP prediction method is seen to offer no performance advantage over the FL predictor on the uncalibrated images. In addition, the table shows that the algorithms that exploit calibration artifacts in 1997 AVIRIS images have no performance advantage on the

³<http://compression.jpl.nasa.gov/hyperspectral>

TABLE IV
COMPRESSION PERFORMANCE (RATE IN BITS/SAMPLE) ON AVIRIS
IMAGES IN THE CCSDS TEST SET.

Image	FL	FL#	LUT#	LAIS- LUT#	TSP- W1	TSP- W2
16-bit Calibrated Images						
Yellowstone 0	3.96	3.91	4.82	4.48	3.94	3.99
Yellowstone 3	3.83	3.79	4.62	4.31	3.81	3.86
Yellowstone 10	3.40	3.37	3.96	3.71	3.37	3.42
Yellowstone 11	3.63	3.59	4.34	4.02	3.60	3.67
Yellowstone 18	3.94	3.90	4.84	4.48	3.92	3.97
Average	3.75	3.71	4.52	4.20	3.73	3.78
16-bit Uncalibrated Images						
Yellowstone 0	6.23	6.20	7.14	6.78	6.23	6.27
Yellowstone 3	6.10	6.07	6.91	6.60	6.09	6.13
Yellowstone 10	5.65	5.60	6.26	6.00	5.59	5.64
Yellowstone 11	5.86	5.81	6.69	6.30	5.83	5.88
Yellowstone 18	6.32	6.26	7.20	6.82	6.28	6.32
Average	6.03	5.99	6.84	6.50	6.01	6.05
12-bit Uncalibrated Images						
Hawaii 1	2.64	2.58	3.26	3.05	2.61	2.62
Maine 10	2.72	2.68	3.45	3.19	2.71	2.74
Average	2.68	2.63	3.35	3.12	2.66	2.68

2006 calibrated AVIRIS images.

Figure 7 shows the band scaling factors for AVIRIS Yellowstone images in the CCSDS test set. Notably, in contrast with the 1997 images, the scaling factor is less than one for all but the first three bands. This is significant because with a scaling factor that is approximately one or smaller we would expect calibration induced artifacts of the sort exhibited in Figure 2 to be significantly diminished or eliminated, as the different DN values would not translate to easily distinguishable peaks in the calibrated sample distribution. In fact we did not discern conspicuous calibration-induced structure in histograms for individual bands of the calibrated 2006 AVIRIS images. As an aside, we remark that scaling a band by a factor less than one followed by rounding to an integer is an irreversible (lossy) operation that reduces dynamic range. Bands from the calibrated images yielded lower compressed data volumes than the corresponding bands in the raw data by an amount very nearly equal to the difference that would be expected from the scaling factors.

VI. CONCLUSIONS

To summarize, the standard 1997 AVIRIS images have some artificial structure that was introduced by radiometric calibration, and compression algorithms can (by design or by accident) exploit this structure. This is the case for the LUT and LAIS-LUT compressors. In fact, we were able to take a compressor that did not exploit these artifacts and make relatively simple modifications (designed specifically to exploit the artifacts) to produce a new lossless compressor that yields roughly a half bit per sample improvement over the previous best (among all compressors) on the 1997 AVIRIS images.

Compressors that exploit calibration artifacts may be desirable for data archival and distribution applications, provided that the data exhibits such artifacts. Furthermore, if calibration

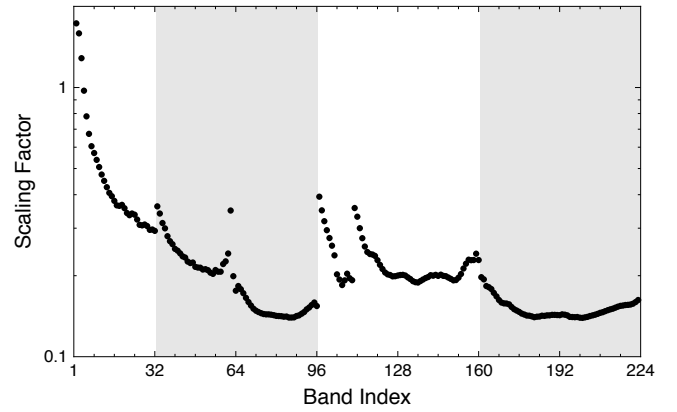


Fig. 7. Band scaling factors between radiance and raw DN values for Yellowstone AVIRIS images in the CCSDS test set.

artifacts are expected in an application, a compressor specifically designed to exploit such artifacts may achieve a large advantage, as evidenced by our TSP compressors.

Our main point, however, is that comparisons among compressors on hyperspectral images that contain calibration artifacts may not translate well to hyperspectral images that do not contain such artifacts. In particular, when the application is compression of raw hyperspectral images onboard a spacecraft, compressor comparisons on the 1997 AVIRIS images can be misleading. In fact, while the TSP compressors and the LUT-based compressors that we examined achieved excellent performance on the 1997 AVIRIS images, they fared poorly compared to the FL compressor on uncalibrated AVIRIS images and even on calibrated AVIRIS images (from 2006) that lack conspicuous calibration artifacts.

We emphasize that there is nothing inherently wrong with a compression algorithm that has the ability to exploit redundancies not seen in raw data. Indeed, theoretical results for universal data compression imply that a compressor can potentially exploit a very wide variety of types of artifacts in an image, while achieving a compression effectiveness only negligibly worse on artifact-free data than a compressor that does not exploit artifacts. However, in practice such a flexible algorithm would likely be more complex (potentially significantly so) than the non-exploiting algorithm, and so would usually not be a serious competitor for raw data applications.

We have illustrated how calibration-induced data structure can be exploited by adaptive entropy coding applied to residuals from a suitably defined predictor, as is the case for the LUT, LAIS-LUT, and TSP compressors. The potential for exploiting such structure is a characteristic inherent in the data set. Other compression methods, such as the use of runlength coding modes, as in [10], and options for entropy coding certain spectral bands directly (rather than encoding prediction residuals), as in [4] and [11], have the potential to benefit from data regularities in more subtle ways. Furthermore, lossy and near-lossless compression techniques can also potentially exploit calibration artifacts; however, we would expect any gains in these cases to be small and to diminish rapidly as more distortion is allowed. It would be prudent to be at least mildly cautious when interpreting any compression results on

the 1997 AVIRIS images.

More importantly, we suggest that compression researchers move away from using the 1997 AVIRIS images as a standard test set. We have made the other AVIRIS scenes mentioned in this article publicly available so that they can be used for compression comparisons. (Note that even though the 2006 *calibrated* images do not contain obvious calibration artifacts, we suggest using the uncalibrated images for compression benchmarking purposes.)

We remark that hyperspectral images are not the only type of data that may contain calibration artifacts; for example, multispectral images and ordinary (two-dimensional) images acquired for science purposes may undergo processing that includes calibration.

It is also perhaps worth noting that some types of detectors may inherently produce data with non-smooth histograms (e.g., some types of analog-to-digital converters can have an odd-versus-even bias). Thus even raw data may contain artificial structure that can be exploited. Clearly, in such a case an onboard data compression system might quite reasonably be designed to exploit such artifacts, although we might expect the gains from such exploitation to be much smaller than those demonstrated here.

More generally, it is important to note that raw data from a spaceborne hyperspectral imager may contain artifacts due to the imager design that are not found in raw AVIRIS images. For example, in a “pushbroom”-type imager, samples in different cross-track positions are acquired with different detector elements that may differ in gains and other properties; as a result, the acquired images may contain features not present in images from a “whisk broom” imager such as AVIRIS. The bottom line is well known: compare data compression algorithms on data that is representative for the desired application.

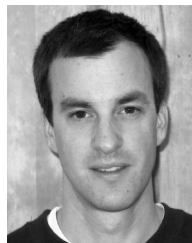
ACKNOWLEDGMENT

The authors would like to thank Dr. R. Green for numerous helpful discussions and for providing AVIRIS data sets; and Dr. F. Rizzo for sharing detailed compression results.

REFERENCES

- [1] J. Mielikainen and P. Toivanen, “Lossless compression of hyperspectral images using a quantized index to lookup tables,” *IEEE Geosci. Remote Sens. Lett.*, vol. 5, no. 3, pp. 474–478, Jul. 2008.
- [2] B. Huang and Y. Sriraja, “Lossless compression of hyperspectral imagery via lookup tables with predictor selection,” *Proc. SPIE*, vol. 6365, 63650L, 2006.
- [3] J. Mielikainen, “Lossless compression of hyperspectral images using lookup tables,” *IEEE Signal Process. Lett.*, vol. 13, no. 3, pp. 157–160, Mar. 2006.
- [4] J. Mielikainen and P. Toivanen, “Clustered DPCM for the lossless compression of hyperspectral images,” *IEEE Trans. Geosci. Remote Sens.*, vol. 41, no. 12, pp. 2943–2946, Dec. 2003.
- [5] B. Aiazzi, L. Alparone, S. Baronti, and C. Lastris, “Crisp and fuzzy adaptive spectral predictions for lossless and near-lossless compression of hyperspectral imagery,” *IEEE Geosci. Remote Sens. Lett.*, vol. 4, no. 4, pp. 532–536, Oct. 2007.
- [6] M. Klimesh, “Low-complexity lossless compression of hyperspectral imagery via adaptive filtering,” *IPN Progress Report*, vol. 42-163, pp. 1–10, Nov. 15, 2005. Available: http://ipnpr.jpl.nasa.gov/progress_report/42-163/163H.pdf
- [7] M. Klimesh, “Low-complexity adaptive lossless compression of hyperspectral imagery,” *Proc. SPIE*, vol. 6300, 63000N, 2006.

- [8] G. Motta, F. Rizzo, and J. A. Storer, “Compression of hyperspectral imagery,” in *Proc. Data Compression Conf.*, Snowbird, UT, Mar. 2003, pp. 333–342.
- [9] F. Rizzo, B. Carpentieri, G. Motta, and J. A. Storer, “Low-complexity lossless compression of hyperspectral imagery via linear prediction,” *IEEE Signal Process. Lett.*, vol. 12, no. 2, pp. 138–141, Feb. 2005.
- [10] H. Wang, S. D. Babacan, and K. Sayood, “Lossless hyperspectral-image compression using context-based conditional average,” *IEEE Trans. Geosci. Remote Sens.*, vol. 45, no. 12, pp. 4187–4193, Dec. 2007.
- [11] S. K. Jain and D. A. Adjeroh, “Edge-based prediction for lossless compression of hyperspectral images,” in *Proc. Data Compression Conf.*, Snowbird, UT, Mar. 2007, pp. 153–162.
- [12] B. Penna, T. Tillo, E. Magli, and G. Olmo, “Transform coding techniques for lossy hyperspectral data compression,” *IEEE Trans. Geosci. Remote Sens.*, vol. 45, no. 5, pp. 1408–1421, May 2007.
- [13] B. Penna, T. Tillo, E. Magli, and G. Olmo, “Hyperspectral image compression employing a model of anomalous pixels,” *IEEE Geosci. Remote Sens. Lett.*, vol. 4, no. 4, pp. 664–668, Oct. 2007.
- [14] Q. Du, W. Zhu, J. E. Fowler, “Anomaly-based JPEG2000 compression of hyperspectral imagery,” *IEEE Geosci. Remote Sens. Lett.*, vol. 5, no. 4, pp. 696–700, Oct. 2008.
- [15] Q. Du and J. E. Fowler, “Hyperspectral image compression Using JPEG2000 and principal component analysis,” *IEEE Geosci. Remote Sens. Lett.*, vol. 4, no. 2, pp. 201–205, Apr. 2007.
- [16] G. Carvajal, B. Penna and E. Magli, “Unified lossy and near-lossless hyperspectral image compression based on JPEG 2000,” *IEEE Geosci. Remote Sens. Lett.*, vol. 5, no. 4, pp. 593–597, Oct. 2008.
- [17] R. O. Green *et al.*, “Imaging spectroscopy and the Airborne Visible/Infrared Imaging Spectrometer (AVIRIS),” *Remote Sens. Environ.*, vol. 65, no. 3, pp. 227–248, 1998.
- [18] R. O. Green, S. A. Larson, and H. I. Novack, “Calibration of AVIRIS digitized data,” in *Proc. 3rd Airborne Visible/Infrared Spectrometer (AVIRIS) Workshop*, JPL Publication 91-28, pp. 109–118, 1991.
- [19] R. O. Green, private communication, Sep. 2008.
- [20] M. J. Weinberger, G. Seroussi, G. Sapiro, “The LOCO-I lossless image compression algorithm: principles and standardization into JPEG-LS,” *IEEE Trans. Image Process.*, vol. 9, no. 8, pp. 1309–1324, Aug. 2000.
- [21] A. Moffat, R. Neal, and I. H. Witten, “Arithmetic coding revisited,” *ACM Trans. Inf. Syst.*, vol. 16, no. 3, pp. 256–294, 1998.
- [22] J. Åberg, Y. M. Shtarkov, and B. J. M. Smeets, “Multialphabet coding with separate alphabet description,” in *Proc. Compression and Complexity of Sequences 1997*, Salerno, Italy, pp. 56–65.
- [23] A. S. Orlitsky, N. P. Santhanam, and J. Zhang, “Universal compression of memoryless sources over unknown alphabets,” *IEEE Trans. Inf. Theory*, vol. 50, no. 7, pp. 1469–1481, 2004.



Aaron Kiely received the B.S. degree in electrical engineering from Virginia Tech in 1989, the M.S.E. and Ph.D. degrees in electrical engineering from the University of Michigan in 1990 and 1993 respectively, and the M.S. degree in aerospace engineering from the University of Southern California in 1999.

Since 1993, he has worked at Caltech’s Jet Propulsion Laboratory. He is interested in spacecraft data compression techniques for a variety of data types including grayscale, multispectral, and hyperspectral imagery; and stereo image pairs on future Mars rovers. He chairs the CCSDS Multispectral and Hyperspectral Data Compression working group.



Matthew Klimesh received B.S.E., M.S.E., and Ph.D. degrees, all in electrical engineering from the University of Michigan in Ann Arbor, in 1989, 1990, and 1995, respectively. He spent one year as a research fellow (postdoc) at Michigan.

Since 1996 he has been with the Information Processing Group at Caltech’s Jet Propulsion Laboratory, working primarily on research and development of data compression algorithms for space applications. His research interests include source coding, data compression, network coding, rate-distortion theory, channel coding, probability, and discrete mathematics.